

Chapter

3

Check Point IPS-1

The Check Point IPS-1 Real-Time Threat Protection product suite includes the most trusted, dynamic and usage-flexible enterprise intrusion prevention system (IPS) on the market today. Based on powerful detection and prevention capabilities, Check Point IPS-1 offers industry-unique features, such as patent-pending Confidence Indexing, customizing capabilities, and minimal design. IPS-1 also supports central management of distributed environments in both small-and-large scale enterprises.

Objectives:

- Identify the Check Point IPS-1 product components and their role in network intrusion protection.
- Describe IPS-1 architectures for common installations.
- Define IPS-1 Passive Mode.
- Define IPS-1 Inline Mode.
- Describe IPS-1 standby options.

Key Terms:

- Packages
- Back-ends
- Network operations center

IPS-1 Components

This section introduces components of the IPS-1 intrusion prevention system:

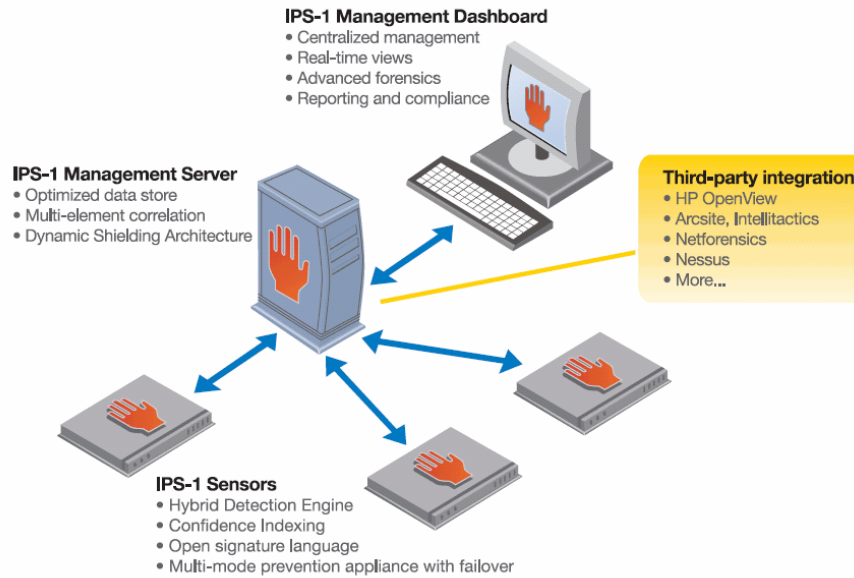


Figure 3-10: IPS-1 Components

IPS-1 Sensor

An IPS-1 Sensor monitors a network for suspicious activity and attacks. These incidents are detected by the **packages** and **back-ends** installed on each sensor. Packages and back-ends contain the actual instructions (N-Code) to filter and process network traffic.

Packages monitor a network for specific categories of exploits. Back-ends monitor the network for specific exploits. For example, the WWW

package has one back-end that watches for attacks against Apache Web servers and one that watches for attacks against IIS Web servers.




Figure 3-11: Network with IPS-1 Sensor

IPS-1 Alerts Concentrator

IPS-1 Alerts Concentrator manages IPS-1 Sensors. It provides central administration and storage of alert and event data for the IPS-1 Sensors that report to it. Multiple IPS-1 Alerts Concentrators can be distributed throughout the network as needed.

When the IPS-1 Sensor detects a possible incident on the network, it generates an alert, which typically includes the name of the package and back-end that identified the incident. **Signatures** are used to detect incidents and cause alerts to be generated. Each signature is then mapped to an alert name. The alert name tells the IPS-1 Alerts

Concentrator which Alert Name, Priority, Description to display in the Alert Browser Window.

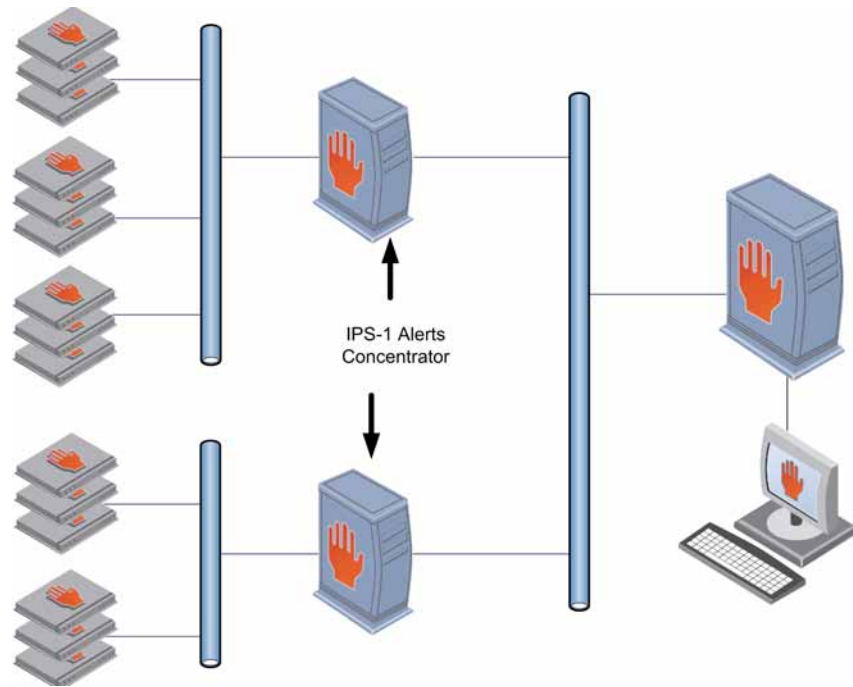


Figure 3-12: IPS-1 Alerts Concentrator

IPS-1 Server

The IPS-1 Server receives alerts from all IPS-1 Alerts Concentrators in the system. At this level, the Security Administrator sets up rules called correlators. These rules tell the IPS-1 Server to take action when it sees a number of alerts that contain identical values within specific fields.

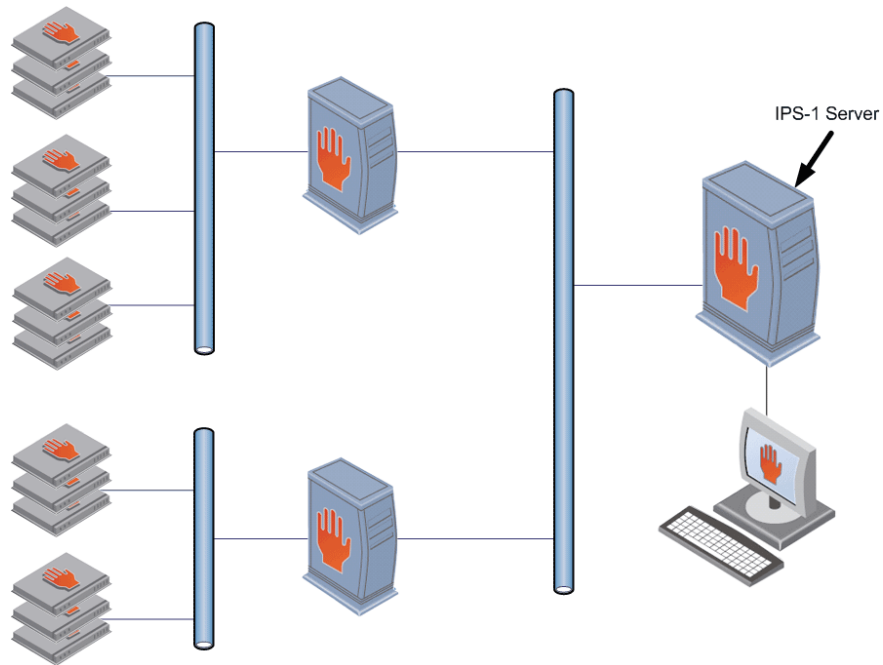


Figure 3-13: IPS-1 Server

IPS-1 Management Dashboard

The IPS-1 Management Dashboard allows you to monitor alerts from the desktop. In addition to monitoring alerts, you can tailor the viewing of alerts according to the network's needs. This tailoring includes viewing alerts by severity, through graphs and time lines, and through the process of selecting alert criteria.

You can also manage IPS-1 components through the IPS-1 Management Dashboard.




Figure 3-14: IPS-1 Management Dashboard

The IPS-1 Management Dashboard is used to view incoming alerts, and:

- To install a new signature.
- To modify the policy on the IPS-1 sensors.
- To receive and manage real-time alerts from the IPS-1 Sensors.

- For forensic analysis of events
- To create graphical reports, etc.

Architecture and Placement

Selecting Placement Points

Placement of sensors is vitally important for a successful IPS deployment. Where should you put IPS devices to maximize their effectiveness? Anywhere your infrastructure or applications are unjustifiably at risk — these areas would be likely targets. Typically, IPS devices are deployed:

- Behind firewalls and WAN routers.
- In front of server farms or similar collections of resources.
- At other network-access points.

By concentrating on these critical points, you will reap the greatest reward from your initial deployment. The reason for this is that most compliance requirements focus on the ingress and egress points to the network core. Also, deploying an IPS at these choke points in the network provides maximum protection opportunities, because they involve transporting and enabling the most network traffic.

WAN router points are excellent candidates for IPS deployment, as they are often the entry points for exploits from remote sites — where you have little direct control and perhaps no authoritative control. If a remote site or business-partner site is compromised, you are often defenseless against an infection already running rampant at that location.

In addition to server farms and other hardened access points, a connection from a wireless warehouse application is another type of access area. Blackberry servers or handheld, wireless barcode readers are examples of this. These areas are especially vulnerable points within any infrastructure. They often mark boundaries within your network, and may represent services and devices that cannot be protected by any other method. These boundaries also represent additional logical and physical responsibilities. These access points represent hard-to-secure applications or services, and must be protected.

IPS-1 Sensor Placement — In Your DMZ

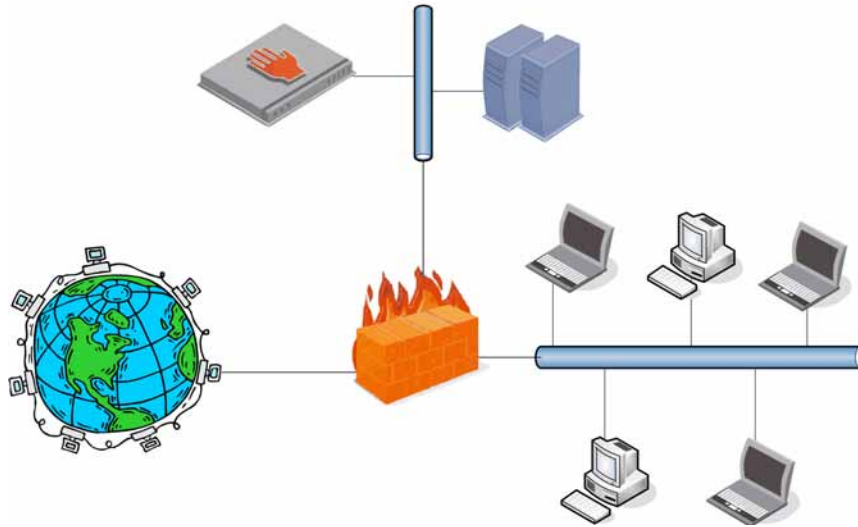


Figure 3-15: IPS-1 Sensor — In DMZ

By placing an IPS-1 Sensor in your DMZ, you are able to monitor any attacks that penetrate your network's perimeter defenses. An IPS-1 Sensor placed here can provide information that will help you improve your secure firewall policies, and watch for attacks that target Web servers and FTP servers. These machines are exposed to the public and are common targets for attack, because if they are breached, they can provide easy access into internal networks.

IPS-1 Sensor Placement — On Network Backbones




Figure 3-16: IPS-1 Sensor — On Network Backbone

By placing an IPS-1 Sensor on major network backbones, you are able to easily monitor large amounts of your network's traffic, looking for any attack that makes it through your firewall. In addition a Sensor placed on network backbones can monitor for unauthorized activity by internal users.

IPS-1 Sensor Placement — On Critical Network Subnets

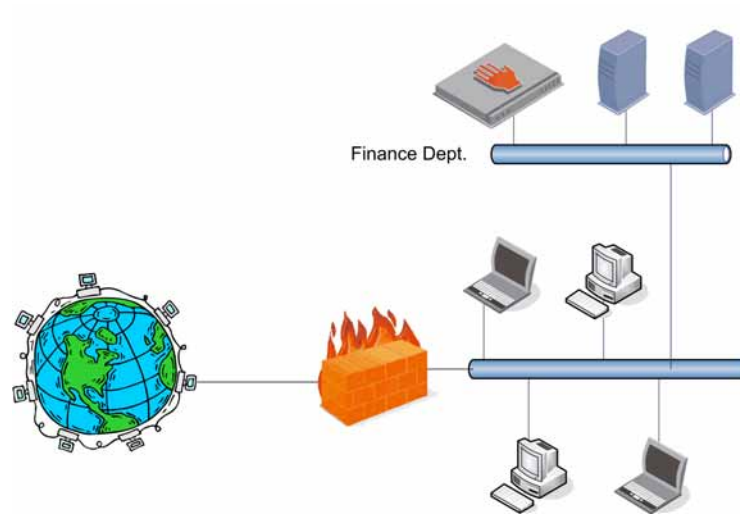


Figure 3-17: IPS-1 Sensor — On Key Subnet

By placing an IPS-1 Sensor on key or critical subnets in your network, you can monitor for attacks that are specifically targeting these critical systems and resources. In a highly switched enterprise network environment, you may want to position a Sensor at each switch, to monitor just that network segment.

Other Possible PS-1 Sensor Locations:

In Front of External Firewalls — Monitoring the number and types of attacks that are targeting your network can help you tune your overall network's security, by letting you know what kinds of attacks you are facing. A Sensor here can also monitor traffic leaving your network, to determine if an internal host has been compromised and taken over.

Remote-Access Servers or Modem Pools — Remote-access points often are not as well protected, and can be exploited by attackers to gain entry into your network.

Wireless access points — Wireless connections into networks are notoriously weak, and are easily discoverable and accessible.

Extranets — Sensors placed at the partner end of an extranet can monitor traffic between you and any key business partners.

Remote Offices — Remote company locations may not have the staff or resources to maintain as high a security level as an enterprise network. IPS-1 Sensors placed in remote offices can monitor for attacks aimed at these less-protected locations.

VPNs — Encrypted VPN traffic cannot not be easily monitored by a sensor, but an IPS-1 Sensor placed on one side or the other of a VPN tunnel will monitor this traffic, either before it is encrypted or after it is decrypted, to help ensure the traffic's integrity.

Deployments

A small deployment has the IPS-1 Alerts Concentrator and IPS-1 Server installed on a shared platform, with IPS-1 Sensors deployed according to need.




Figure 3-18: Basic Deployment

However, large enterprise deployments will need to distribute the data and processing required to handle large volumes of data across multiple servers. In this case, the applications running on the Protection Engine are distributed. At the top level is the IPS-1 Server, while IPS-1 Alerts

Concentrators can be broken out separately. You may have as many Alerts Concentrators as you like communicating with an IPS-1 Server.

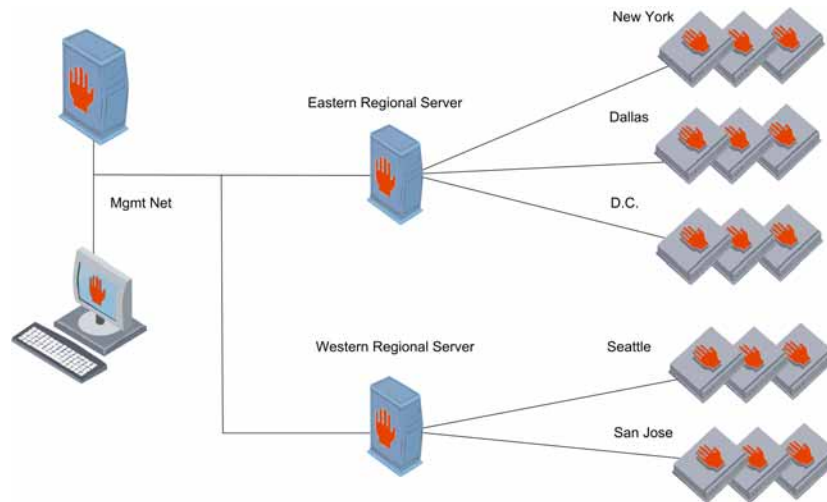


Figure 3-19: Enterprise Deployment

The IPS-1 Server is a central alert and configuration repository. IPS-1 Sensors can record a lot of data, for example, data recorded when a user tries to log in to an FTP server as root. First, a single alert is generated, which creates one small record. Next assume packet capture is turned on with the default of 10 packets, which generates 10 more records. So, one alert generates 11 records, 10 of which could be quite sizeable, since it's raw packet capture information. And as previously mentioned IPS-1 Sensors can do more than record alerts, as shown with general web and tcp sessions, and others. You might need to have multiple IPS-1 Alerts Concentrators just to keep up with the rate at which your Sensors are spooling data.

Rather than try to move around all these records, the IPS-1 Server simply receives the Alert Record. It leaves the other records on the Alerts Concentrator. When a user wants to retrieve those records, the IPS-1 Server retrieves them from the appropriate Alerts Concentrator.

This reduces the need to duplicate data, and also the need to move large amounts of data around the network until it is actually needed.

So, for geographical reasons you might have multiple IPS-1 Alerts Concentrators, especially in a global system, where you might have groups of IPS-1 Sensors separated by continents. There is no need to move large amounts of data, if you are not always going to use it.

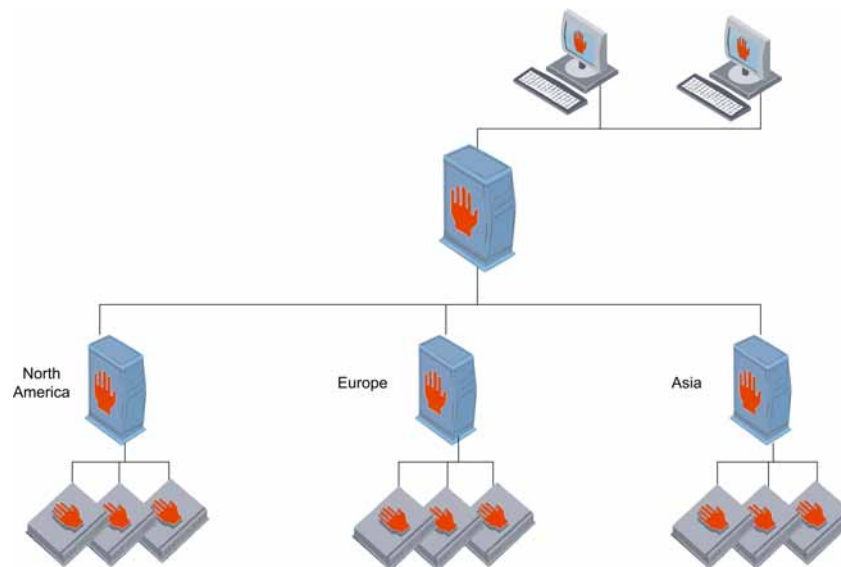


Figure 3-20: Geographic Deployment

Organizational Deployment

There may be political reasons to have multiple Alerts Concentrators. For example, a large organization may standardize on a product like IPS-1, so each division of the company might be using the product, for example, the East Coast and West Coast division of a consulting company. Or a large manufacturing organization might use IPS-1, with one product group fairly isolated from the other groups. Within that organization, each group may want to manage their own systems, to stay autonomous. But the global organization may only have funds for one 24x7 **Network Operation Center (NOC)**. So, each sub organization controls its own IPS-1 Alerts Concentrator, but the 24x7 NOC has the IPS-1 Server.

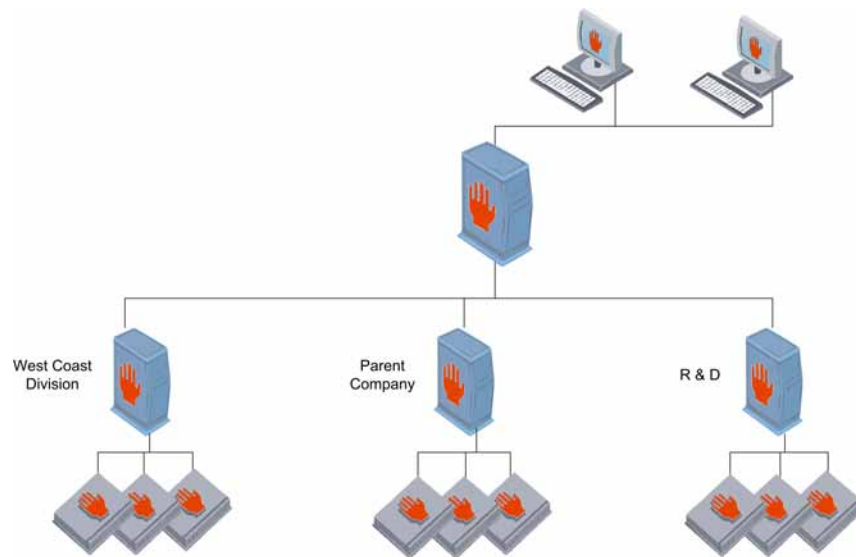


Figure 3-21: Organizational Deployment

High Availability

IPS-1 provides alert continuity through multiple High Availability Servers.

This configuration option ensures continuity of information from IPS-1 Sensors in the event of an IPS-1 Alerts Concentrator failure. You can configure IPS-1 Sensors to report to a backup IPS-1 Alerts Concentrator. This automatically redirects alerts and event data to the backup IPS-1 Alerts Concentrator if the primary IPS-1 Alerts Concentrator fails. You can install the backup IPS-1 Alerts Concentrator within the same network as the primary IPS-1 Alerts Concentrator.

As shown in the following diagram, if there are multiple IPS-1 Sensors within your enterprise, you can designate one sensor's primary IPS-1 Alerts Concentrator as the backup IPS-1 Alerts Concentrator for another Sensor.

With the IPS-1 Management Dashboard, if an Alerts Concentrator fails, you will not see the alerts unless you are already connected to the Secondary Alerts Concentrator. But if you're using the IPS-1 Server, it does not matter where the alerts come from. This is because you are using the same Management Dashboard for both Alerts Concentrators.

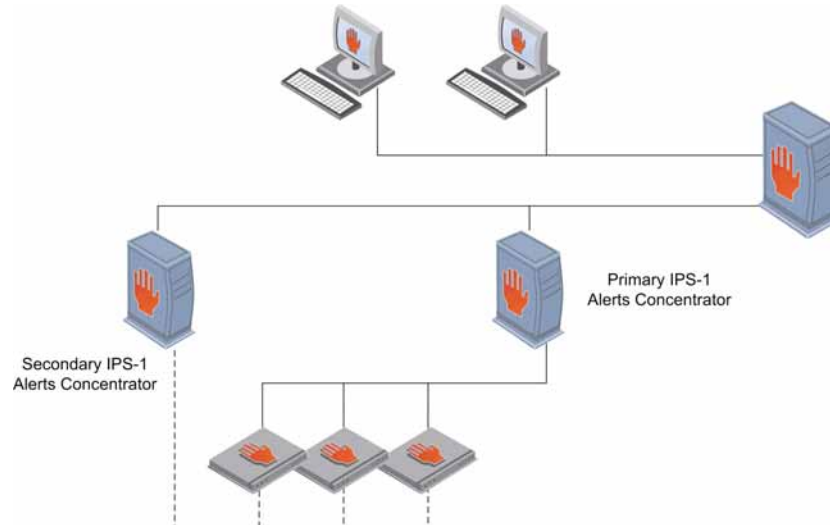


Figure 3-22: High Availability

IPS-1 Sensor Installation Modes

IPS-1 may be deployed in four different modes: Passive Mode, Inline Bridge Mode (also known as Learning Mode), Inline Fail-severed Mode, and Inline Fail-passthrough Mode.

Passive Mode — As the traditional intrusion detection mode, Passive Mode receives traffic from a span port or tap. In this mode, IPS-1 is not inline but receives a copy of the traffic for attack-detection analysis, and will raise alerts accordingly. Each non-management interface is available to monitor a span port or hub.

Inline Bridge Mode — When deployed in Inline Bridge Mode (also known as Learning Mode), IPS is deployed inline but it does not actually prevent attack traffic when detected. In this mode, the sensor will never discard or drop a packet; it acts as a repeater to the network and a pass through Intrusion Detection device to the Administrator. This mode enables a transition step from detection to prevention, where the user may observe the effectiveness of IPS-1 running inline, detecting attacks, and without actually dropping any detected attack traffic. This mode is also useful for the network management team to assess the device's resiliency for placing it in full Inline Prevention Mode. A copy of each packet is made, this copy is then checked by the Sensor, and acted upon if necessary. If not acted upon, the packet copy is then "black holed". This process helps limit any network performance degradation through the Sensor

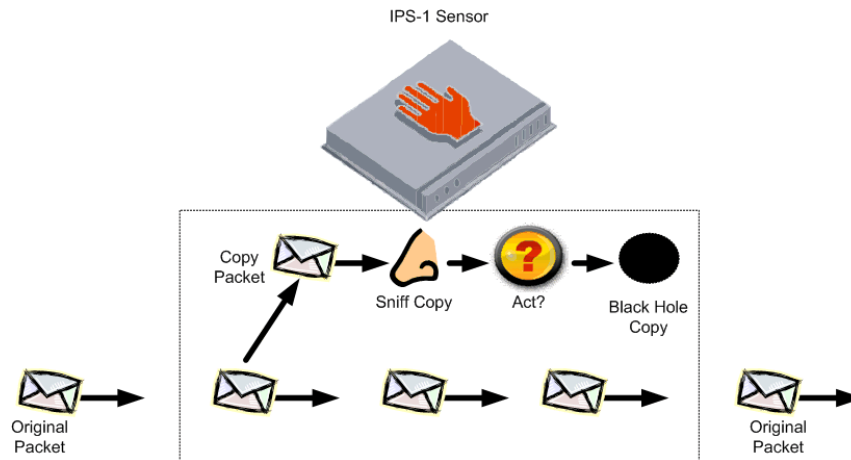


Figure 3-23: Inline Bridge Mode Packet Checking

Inline Fail-passthrough Mode — In the fail passthrough mode, if the Sensor fails, network traffic will continue to pass through, ensuring network integrity and business operations. This typical default mode ensures continued operations of the network with no impediment or loss of traffic.

Inline Fail-severed Mode — The fail severed mode results in no pass-through of network traffic if a sensor fails. Using this mode enforces a policy where security is more important than traffic flow.

IPS-1 Sensor Inline-Mode Packet Flow

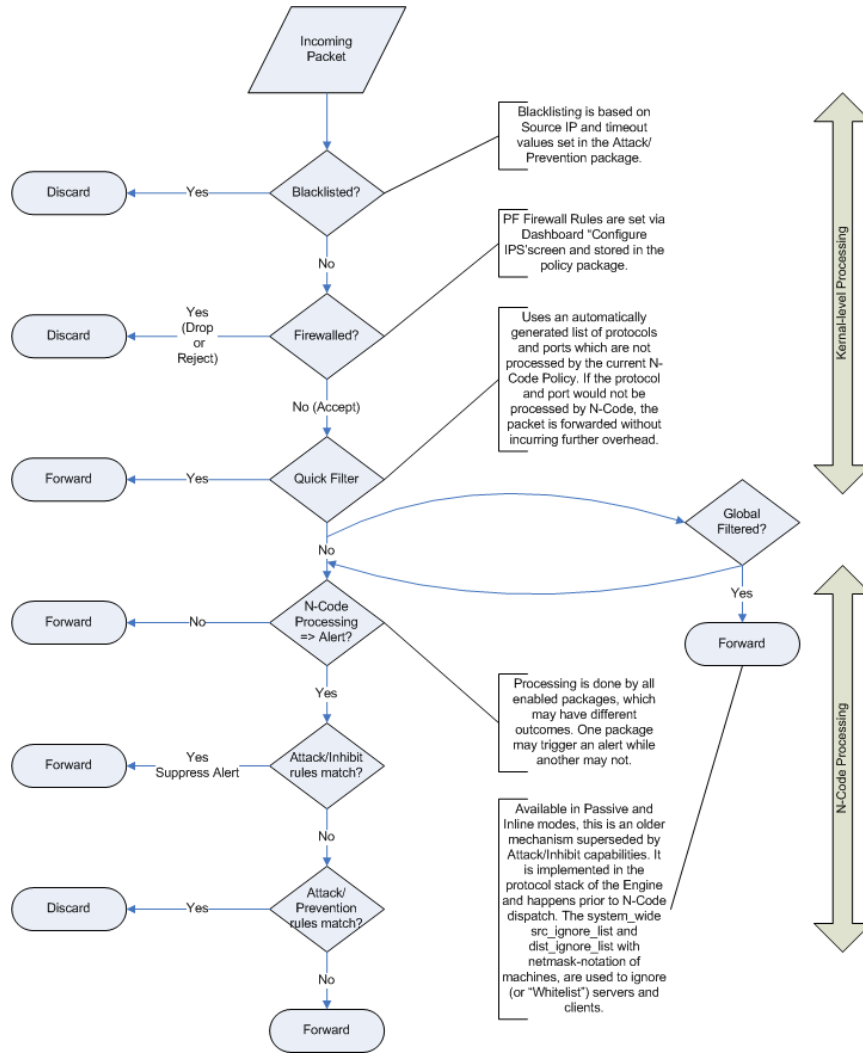


Figure 3-24: Inline-Mode Packet Flow

